2008 IEEE International Conference on
Robotics and Automation
Pasadena, CA, USA, May 19-23, 2008

An Extremal Fields Approach for the Analysis of Human Planning and Control Performance

Bernard Mettler
Department of Aerospace Engineering and Mechanics
University of Minnesota
Minneapolis, MN 55455

Abstract— **In this paper we describe a study of the human pilot control behavior in a planar goal directed flight task. The experimental data was collected using a miniature helicopter in an indoor flight test facility. To provide insight into the human's control behavior we developed a technique to extract extremal fields from the family of collected trajectories. These fields describe the spatial distribution of the vehicle states and cost-to-go, including their statistical distribution, which provides information about the variability of the pilot's control behavior over the task domain. Once extracted we can compare these fields to the value functions obtained from the task's equivalent optimal control problem. The comparison of the human-extracted and the computed value function maps suggests that on average, the human acts similarity to an optimal control policy. The results also suggests that a simple mass-point model used for our analysis, and motivated by the hypothesis that the pilot acts as a dynamic inverse controller, is sufficient to explain the pilot's performance at the planning level. We use these results to develop hypotheses about human planning and control processes and discuss their biological plausibility based on control-theoretic interpretations. We plan to use the new insights from this framework to help design more capable and versatile algorithms for autonomous vehicle control, as well as help design man-machine interfaces that enable a more natural link with the operator's internal control and planning processes.**

I. INTRODUCTION

Human pilots excel in controlling vehicles with complex dynamics like helicopters. Trained pilot can seamlessly execute maneuvers from a broad repertoire to negotiate challenging three dimensional environments under often disturbed and uncertain conditions. These skills speak for tightly integrated planning and control processes; the pilot essentially seems to incorporate the vehicle into his own spatial behavioral unit.

Trajectory planning is a classic problem in robotics and control theory, and is one of the key capability needed for autonomous vehicles. For agile vehicles, existing methodologies are still far from enabling performance levels akin to human pilots. We know little about the exact principles and processes underlying of human control skills. Piloting is an intensive form of exercise for the brain's control, perceptual, and cognitive processes therefore human flight experiments are a rich source of psycho-physical information. Human pilots have mostly been studied in relation to issues of handling qualities, yielding models that describe the pilot's role as a regulator [1]. More recently pilot's control strategies during acrobatic maneuvers have been studied [2]. Simple maneuvers in highly nonlinear regime were successfully modeled and then reproduced experimentally using hybrid control schemes [3]. Studying the human pilot in tasks that involve planning as well as control could provide insight into fundamental questions of implementation and integration of both the planning and control processes and, at the same time, could help better understand the biological underpinnings of human piloting skills.

In the following, we describe a framework to analyze the human planning and control behavior from an optimal control standpoint. We use a miniature remote-controlled (RC) helicopter in a specially designed indoor flight test facility. To exercise the pilot's control skills, we use a simple yet challenging goal-directed control task. The problem is characterized by the nonlinear vehicle dynamics and the multi-dimensional trajectory state-space (x, y planar coordinates, heading ψ and velocity v). To provide insight into the human's planning and control behavior we performed the experiment over a range of initial conditions and repeated each experiment several times to provide sufficient statistical significance. To have a basis to evaluate the pilot's performance we use the solution for the corresponding time optimal control problem. Instead of comparing the trajectory data directly, we introduce a framework based on spatial distribution of the states and costs. These maps can be related to the concept of extremal fields and value functions providing an insight into the control behavior of the human operator from an optimal control standpoint.

II. BACKGROUND AND MOTIVATION

Model-based, trajectory optimization, provides a rigorous mathematical approach to formulate the combined planning and control problem. However, optimal control techniques are computationally expensive and also provide no natural way to account for a geographical environments in its global form. Historically, the theory and technologies have been developed to tackle aerospace problems where the environment is trivial and largely known and the conditions are relatively stationary. The typical applications include: rocket guidance, spacecraft orbit transfers, and airplane automatic landing [4], [5]. Highly interactive guidance tasks in partially known and even dynamic environments require to updating the trajectory continuously in real-time to account for the

978-1-4244-1647-9/08/$25.00 ©2008 IEEE.

latest information about the environment, the current vehicle state and compensate for disturbances and uncertainties in vehicle behavior.

Recent work has focused on developing computationally more efficient trajectory optimization techniques that can be implemented in real time. One concept is to break down the trajectory optimization problem into an online and an offline portion. The framework of model predictive control (MPC) implemented in receding horizon (RH) fashion provides a formal way to implement such a control scheme [6]. It has been explored for controlling nonlinear systems [7], [8], [9], and more recently has been applied to online trajectory optimization [10], [11], [12]. In this approach, the trajectory is optimized online over a finite time-horizon, based on a prediction model and the current vehicle state. The computation is repeated as the vehicle progresses through the environment. To ensure stability of this control scheme, the discarded tail of the trajectory has to be accounted for by its cost to go. For sake of performance (in the sense of optimality), the cost-to-go should be a good approximation of the actual cost of the discarded portion of the trajectory, i.e., the cost-to-go should be an approximation of the optimal value function (VF). For the computation of a cost-to-go function, heuristic techniques based on geometry [10] and cell decomposition [12] (akin to Navigation Functions) have been used. Recently more exact solutions based on approximations of the spatial value function (SVF) have been explored [13]. The accuracy of these SVF directly influences the agility with which a vehicle can negotiate complex spatial environments.

A. Human spatial control behavior

The development of approximation techniques for trajectory optimization raises questions about how the brain performs spatial control tasks and if optimization principles play a role in these processes. Optimality has been proposed as a fundamental principle in movement planning, control, and estimation [14], [15]. Optimal feedback control principles have been used to explain the coordination of movements and the tradeoff between movement variability and task performance [16]. However, little is known about how optimization principles may play a role in more complex spatial behaviors.

The research in spatial behavior goes back to O'Keefe's work demonstrating that spatial information is encoded by networks of neurons in the hypocampus (so-called place cells) [17]. Such neuronal maps are believed to be used for navigation and spatial reasoning and decision making. Recently, so-called grid cells have been identified that are believed to act as grid-like reference system that would support the spatial mapping process [18]. These new findings provide a more detailed picture of how spatial information is learned and encoded, however, little is known about how these maps may be used in complex spatial control tasks. Simple, goal-directed navigation can be explained using theories of reinforcement learning [19] and computational models based on these maps have been implemented on a simple mobile robot [20]. However, these results are far from reproducing the spatial control skills demonstrated by humans and other animals.

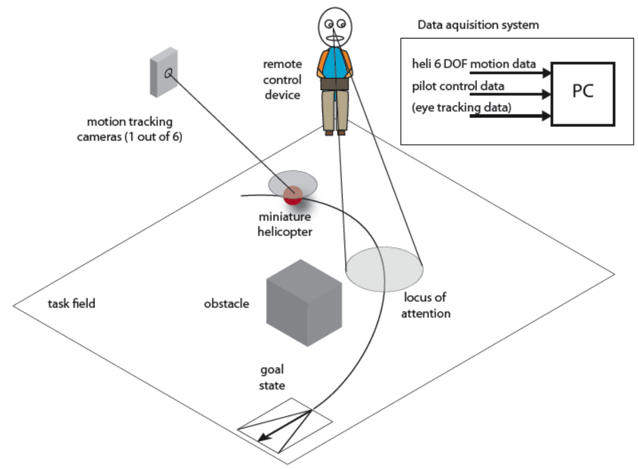

Fig. 1. Illustration of the experimental setup based on our indoor flight test facility. The pilot control inputs are recorded simultaneously to the vehicle 6-DOF data. (The eye tracking device will be used in future experiments.)

III. Experimental Setup

We use a control-theoretic analysis of human spatial control performance to provide clues about how different processes (spatial representation, planning, and online control) are integrated and implemented. To implement and evaluate planning and control algorithms as well as to perform experiments with the human operator exercising his spatial control skills we use an indoor flight experimental facility [21]. The facility uses miniature RC helicopters which can operate in all three dimensions and can exhibit a broad range of dynamic behaviors [22]. These attributes make them a unique tool to collect psychomotor data. When operated by a human pilot, only the fingers are involved minimizing the involvement of the typical bio-mechanical component. Finally, we designed experiments that can be easily formulated as an optimal control problems.

A. Experimental facility

The experimental facility consists of: (1) the RC helicopter; (2) a motion tracking system; (3) an input device (RC transmitter); and (4) the data acquisition system. The RC helicopter in our experiment is a small off-the-shelf Blade CX (34.5 cm rotor diameter and 200 g). It is controlled by a standard four-channel radio transmitter. Thanks to the vision based motion tracking system the required on-board electronics is limited to the standard 4 axis receiver.

The motion capture system consists of 6 high-resolution MX-40 gray-scale cameras from Vicon Systems. We use the ViconIQ software running on a PC to process the information received from the cameras and generate the 6 degrees of freedom (DOF) motion tracking data. The four pilot inputs are recorded simultaneously and integrated into a single data stream.

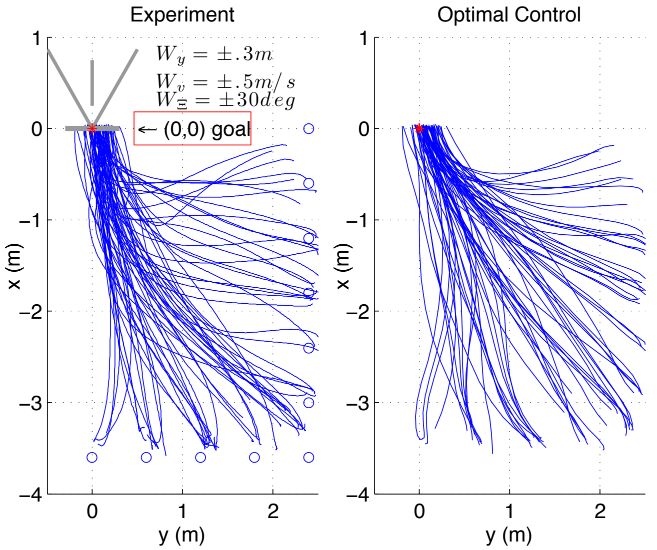

Fig. 2. (left) Trajectory segments from the experiment. The subject has to fly the helicopter from 11 starting locations (shown by the circles) to the goal state situated at the origin (0,0) and specified by a lateral position offset W_y, a velocity magnitude tolerance W_v, and a heading tolerance W_ψ). (right) Trajectory segments computed based on the equivalent optimal control problem.

Retro-reflective spherical markers are placed on the vehicle. The geometrical configuration of the markers is provided to the tracking software which then treats it as a rigid body. The markers tracking accuracy is about 0.02 pixels which corresponds to about 40 μm in our experimental space of 7x7 meters. This level of resolution can be achieved at a sampling frequency up to about 100 Hz. The infrastructure is described in more details in [21].

B. Spatial control experiment

For this first study, we defined a planar goal directed guidance task. For this task illustrated in Figure 2, the pilot has to control the vehicle from multiple initial positions arranged along the edge of a rectangular space to a goal state. The goal state is given by: coordinate $\mathbf{x}_{goal} = (0, 0)$, velocity $v_{goal} = 1$m/sec and heading $\psi = 0$deg.

Tolerances were used to define an acceptable range of variation (the goal set): lateral position tolerance ($W_y = \pm.3$ meters), final speed magnitude tolerance ($W_v = \pm 0.5$ m/sec), and final heading tolerance ($W_\psi = \pm 30$ deg). These values were chosen based on the skills of our pilot at performing this particular task.

For the initial state, the pilot had to start the trajectory from a stationary hover above one of the specified initial positions (within a circle of .25m radius). The initial heading was left open. Since there is significant variability in the human performance, several recordings were made for each starting location.

C. Processing and selection of the trajectories

Following the data collection, the trajectory data was processed to extract the set of trajectory segments used for the analysis. This selection process involves removing faulty trajectories (in particular data losses) and trajectories that do not satisfy the task specifications (outside the goal state region). The data was processed automatically using our segmentation script.

IV. Optimal Control Framework

The family of optimal trajectories going from different initial points to the goal set is known as a field of extremal in calculus of variations [4]. This field describes the control behavior over a region of the planning space. Associated with each point in this field is a value of the performance objective minimized by the optimal control action: the optimal value function $V(x)^*$. From this perspective, our collected trajectories represent the field of extremals from the human operator. We can evaluate the performance of the human operator by comparing its corresponding field with that of the equivalent optimal control problem.

A. Optimal Value Function

Given an optimal control problem described by a performance objective $J(x)$ and a goal state x_{goal}, the optimal value function $V^*(x)$ gives the optimal cost to reach the goal (cost-to-go) from a given state x.

The value function is the solution of the Hamilton-Jacobi-Bellman (HJB) differential equation; a solution to the HJB differential equation represents a sufficient condition for optimality [23].

Most problems of practical interests, however, have no analytic solutions to the HJB equations. The discrete-time equivalent, the functional equation of dynamic programming (DP), is typically used in engineering problems [24].

Since the optimal value function V^* fully characterizes the control problem (including the system dynamics and the performance objective) we can regard $V^*(x)$ as the information needed to perform a particular control task. Hence it is natural to ask whether the human brain encodes something akin to a Value function.

B. Optimal control trajectories

To have a baseline for our analysis we compute the value function for the equivalent optimal control problem corresponding to the task. Since there is no analytical solution to the HJB equation we compute the optimal trajectories (field of extremals) for the start-end conditions obtained from our experimental data. We will then be able to extract a value function for both the experimental data and the computed, optimization based data using the same cell averaging technique.

1) Helicopter model: At the level of the spatial control task, the only variables that are needed explicitly to describe the behavior are the inertial frame positions x, y, the body velocity v, and course angle ψ. Yet the helicopter has complex dynamics which require the coordination of four physical control inputs $u_p = [\delta_{lon}, \delta_{lat}, \delta_{ped}, \delta_{col}]'$. Its dynamic behavior is described by at least a 10th order system (three translational velocities, Euler angles and rates, and the lateral and longitudinal rotor flapping states [22]).

For a skilled pilot it is reasonable to assume that the pilot effectively acts as a dynamic inverse control law. This implies that at the spatial planning level, the pilot abstracts out details of the vehicle control. To compute optimal trajectories, we use the dynamics of a simple mass point model:

$$\begin{aligned} \ddot{x} &= a\cos(\Xi + \theta) \\ \ddot{y} &= a\sin(\Xi + \theta). \end{aligned} \tag{1}$$

The control inputs are the acceleration magnitude a (which can be negative when decelerating), and θ which represents the angle between the current acceleration and course angle Ξ.

The control performance of the pilot as a dynamic inverse controller is captured by constraints on the acceleration and velocity. The numerical values are based on actual flight test experiments. The longitudinal and lateral acceleration limit are given by the constraints: $|a| \leq 1.5m/sec^2$,$|\theta| \leq 30deg$), and the limit on velocity magnitude ($\sqrt{\dot{x}^2 + \dot{y}^2} \leq v_{max} = 2.5m/sec$).

2) Optimal control problem formulation: The optimal control problem for the trajectory planning task is formulated as:

$$\begin{aligned} \min_{\mathbf{u}} J(x) &= \int_0^{t_f} dt \\ \text{subject to} & \\ \dot{\mathbf{x}} &= \mathbf{f}(\mathbf{x}, \mathbf{u}) \\ \mathbf{x}(t_0) &= \mathbf{x}_0 \\ \mathbf{x}(t_f) &= \mathbf{x}_{goal} \\ \mathbf{u} &\in U \\ \mathbf{x} &\in X \end{aligned} \tag{2}$$

where $\mathbf{f}$ is the vector differential equation representing the vehicles equations of motion and U and X are the sets of admissible controls and states, respectively, the acceleration and velocity limits.

We generate two sets of trajectories. The first one assumes fixed end state and optimal initial state for positions distributed evenly along the starting edge. The second set uses the actual start/end states from the experiment. This enables us to capture some of the variability inherent to the human motor and sensing system. For example the lateral arrival position has a mean E(y) = 0.066 m and standard deviation σ_y =0.092 m (notice a slight shift to the right also visible in the Figure 2). The optimal trajectories were computed using an interior point numerical optimization tool [25].

C. Extracting value function and extremal maps

To extract the extremal fields from the family of experimental and computed (optimal control) trajectories, we perform a spatial averaging. We discretize the $x - y$ plane using a regular grid with resolution d_{xy}. For each spatial grid point (i, j), we average the states and cost-to-go values for all trajectory time samples that fall within a circular region of radius r_{xy}. The set of time samples $K_{i,j}(n)$ for a grid point (i, j) and trajectory n in the family is defined by:

$$K_{i,j}(n) = \{k | d_{i,j}(n) \leq r_{xy}\} \tag{3}$$

$$d_{i,j}(n,k) = \|(x_n(kT_s), y_n(kT_s)) - (x_i, y_j)\| \tag{4}$$

The field's values are computed using a weighted average where the weights are the normalized distances from the grid point. For the velocity magnitude, we have:

$$\bar{v}_{i,j} = [\sum_{K_{i,j}(n);n=1:N} w(n,k)_{i,j} v(kT_s)] / \sum_{K_{i,j}(n);n=1:N} w(n,k)_{i,j} \tag{5}$$

$$w(n,k)_{i,j} = [1 - d_{i,j}(n,k)/r_{xy}] \tag{6}$$

For our study, the same averaging process is performed for the course angle $\Xi(i, j)$ and the cost-to-go value $T(i, j)$. The cost function in our interception task is travel duration to the goal set. We evaluate the cost-to-go by determining the time instant for each trajectory n when it reaches the goal state $t_{arr}(n)$ and obtain the CTG value based on the time at the current instant $T = kT - t_f$.

When averaging we also compute the standard deviations (σ_T, σ_v, σ_Ξ). These statistics capture the variability in these quantities which are important for the interpretation of the results.

Figure 3 shows the field values that were extracted through spatial averaging for the experimental and the optimal computed trajectories. Figure 4 shows the corresponding statistics. These results were computed for cells with edges d_{xy} = 0.1m and circular averaging region r_{xy} = 0.6m. These settings provide sufficient spatial resolution and smoothness based for our current number of trajectory segments, and based on the sampling interval $T_s = 0.02$sec.

V. Results

From the extracted fields shown in Figure 3 we see that the human planning and control behavior is consistent enough as a function of space to result in smooth spatial distributions both for the cost and for the state variables. Also, the cell resolution is sufficient to capture the overall shape of the distribution. By increasing the cell size we would loose details about the shape of the distribution.

The speed distribution $\hat{v}$ exhibits the typical bell-shaped profile along the x- direction. Overall, the pilot is operating the vehicle at a slower speed than the computed optimal value. In particular there is a pronounced indentation in speed left of the field's diagonal. This area happens to be where the pilot turns the vehicle toward the goal (see Figure 2). In contrast, in the computed data the heading changes more gradually as it moves toward the goal. This behavior is visible in the heading distribution, where the average heading is relatively flat compared to the computed data.

The time-to-go for the pilot data is larger than for the optimal computed data. The difference of up to 1 sec is not surprising given that the parameters (the acceleration and speed constraints) used for the optimal control problem

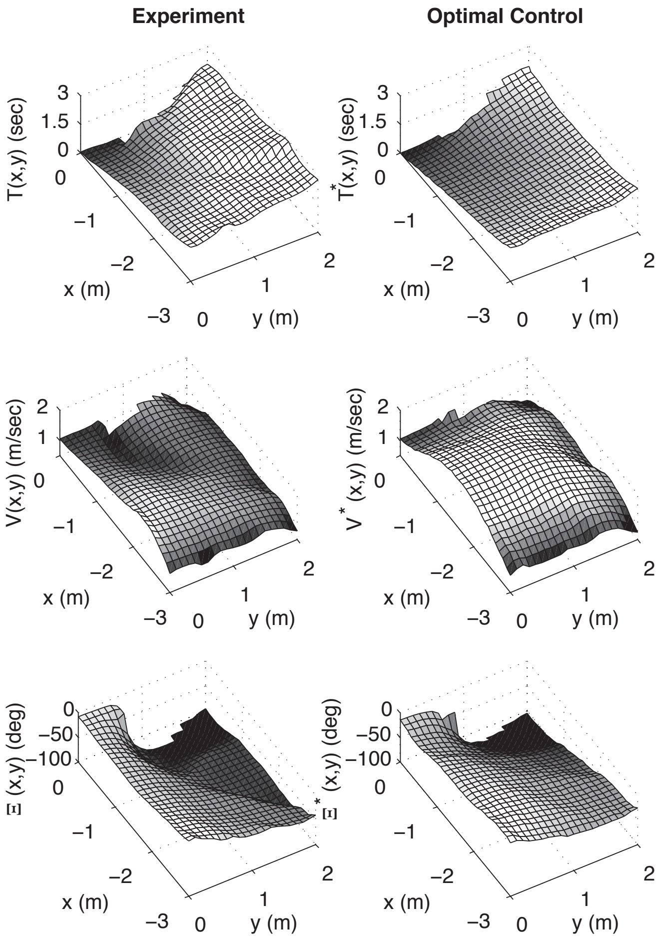

Fig. 3. Spatial distribution of the values extracted from the experimental (left) and computed trajectories (right): the cost-to-go T (time to reach the goal); the speed of the vehicle V; and the course angle Ξ of the vehicle. (The goal is at the upper left corner (0,0))

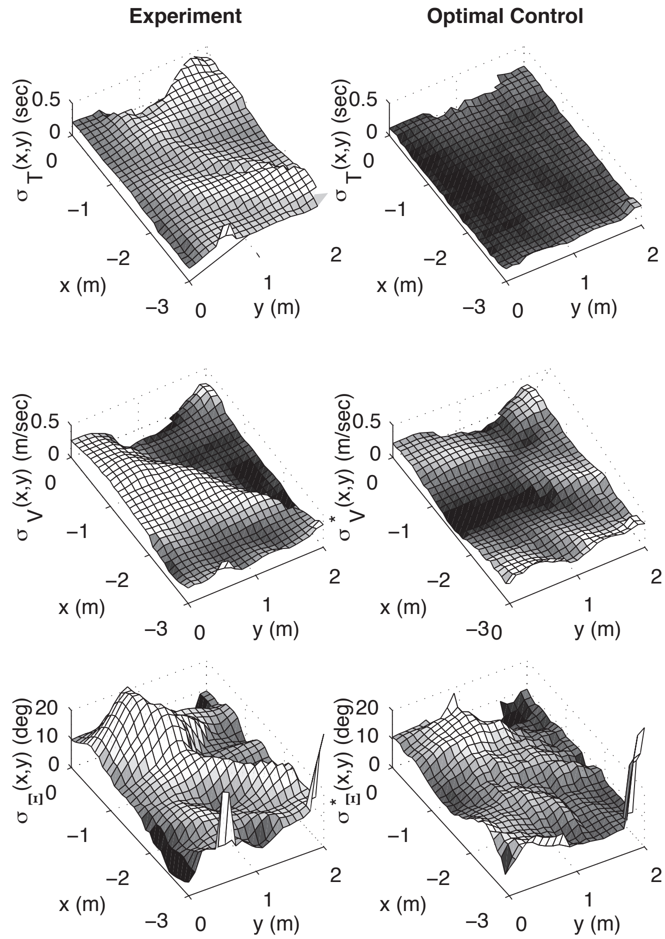

Fig. 4. Spatial distribution of the standard deviation for the experimental data (left) and computed trajectories (right): the cost-to-go σ_T (time to reach the goal); the speed of the vehicle σ_V; and the course angle σ_Ξ of the vehicle.

are based on the particular pilot's best performance. The computed data essentially represents a performance upper bound based on the current skills of the pilot. We could have easily generated an absolute upper bound (i.e. cost-to-go $T(x,y)$ lower bound) based on the vehicle physical limits. This upper bound would indicate how well the pilot performs based on the theoretical performance limit (note that these do not account for the the human sensory and motor control limitations).

The variability of the extracted values is described by the standard deviation and shown in Figure 4. As expected, the experimental data has more variability than the computed data. The highest variability occurs on the task domain's diagonal for the course angle Ξ (up to 20 degrees) and on the left side of it for the velocity v. I this area the goal state is offset and rotated relative to the vehicle motion. From a control perspective this area requires more coordination between the longitudinal and lateral control inputs, thus the performance will depend on the pilot's skill level.

VI. Control Theoretic Interpretation

The field of extremal extracted from the experimental trajectories describe the operator's "closed-loop" performance. These fields, however, do not tell us how the human generates the trajectories and implements the necessary control actions.

From a control theoretic standpoint, at least three different forms of implementations could be used for such spatial control problems. Each form of implementation has implications that can help us evaluate their biological plausibility. All three forms are based on the assumption that at the "physical" control level the human acts as a dynamic inverse controller. Figure 5 illustrates this model. The vehicle dynamics are described by P and P^{-1} is the dynamic inverse control law. Based on current estimates of the vehicle position and state $\hat{x} = [\hat{x}, \hat{y}, \hat{\psi}, \hat{v}]'$ (obtained via the human sensory system), the pilot generates the appropriate physical inputs u_p based on the difference between the optimal state $x^*(x) = [v^*, \psi^*]'$ provided by the planning level policy and the current estimated state.

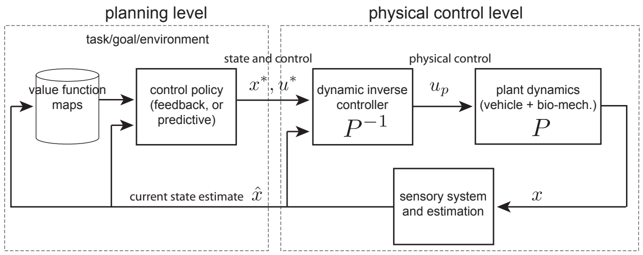

Fig. 5. Idealized model of the human planning and control process. The model assumes that operator controls the vehicle using a dynamic inversion: given the desired vehicle velocity and course angle, the human determines the system's physical inputs. At the planning level different strategies are plausible: 1) in feedback policies, value function maps provide information about the optimal heading and velocity based on the current vehicle state; 2) in full horizon optimization, a trajectory from the current state to the goal is computed at once; 3) in receding horizon predictive optimization, a local trajectory is computed based on an approximation of the value function map.

A. Possible Control Models

1) State Feedback policy: The fields are used as spatial feedback policy maps. The maps provide information on the optimal vehicle state values $x^*(x) = [v^*, \psi^*]'$, i.e. what course and velocity the vehicle should be following, given its current location x, y. The value function $V^*(x)$ provides the time-to-go from that location to the goal, provided the optimal policy were implemented all along the future trajectory.

Such a feedback policy would require that the entire map is available to to the human (as a form of knowledge). While possible, it is not very realistic as a general strategy considering its specificity to a particular task and conditions. However, if a particular skill is used extensively, and the different parameters that affect the map's validity (dynamics, end-state, disturbances) do not change much it may be a plausible model for the control process.

2) Full-horizon optimization: In the full-horizon optimization, a trajectory and corresponding control history is "computed" at once (from scratch) based on current state, task, and a model of the close-loop vehicle behavior. This *nominal* trajectory can then be fed open-loop. From an optimal control viewpoint this is a two-point boundary value problem (BVP) which can be computationally expensive to solve. In the event of large disturbances or significant model miss-match, this nominal trajectory would not be valid and a new solution would have to be computed.

Such a solution implies that the brain performs a large optimization problem online. Such a process is plausible for situations where the subject can utilize much of its "online" processing resources (e.g. when training). However, in challenging conditions, in the presence of significant uncertainties (about the task or environment), the process would have to be repeated almost continuously to update the trajectory. In this case, the necessary "online" processing may exceed the available resources.

3) Receding-horizon predictive optimization: In receding horizon (RH) optimization, only a local trajectory need to be computed (based on current state, task and a closed-loop model) for a finite prediction horizon. To preserve optimality and stability, the discarded tail of the trajectory has to be accounted for. One way of achieving this is by using a cost-to-go function that captures the cost of the discarded tail. The lighter online computation required allows to frequently update the trajectory and hence account for evolving knowledge of the environment and task as well as compensate for exogenous effects.

From a biological standpoint, this model is the most general and provides the most versatility to deal with uncertain and changing tasks and conditions. In fact, for long prediction horizons and accurate internal closed-loop models, the performance of the RH scheme should converge to that of a full-horizon optimization. Conversely, for very short horizons, the scheme becomes more like a state-feedback policy.

VII. Conclusions and Future Work

The extracted extremal fields show that the human performance exhibits good spatial regularity. Further, the comparison with optimal control data suggests that for our particular task the human pilot behaves similarly to a minimum-time, optimal control policy. Moreover it shows that an idealized mass-point model, motivated by dynamic inverse control, is sufficient to explain the pilot's control behavior at the planning level. More generally, the results demonstrate that the concept of extremal fields provides unique insight into the control behavior as a function of space. This is key for the type of spatial control tasks which are common to many human or animal activities. In addition, this framework also provides a spatial perspective on the statistical information. For human control behavior, variability is a typical measure of skill level [26]. The extremal field technique may be a way to extend these notions by relating skill level to a measure of optimality.

The value function is also representative of the information needed to perform the task. Therefore, it may provide a framework to study the information-theoretic aspects underlying human control performance. At the control level, it should help us understand the type and amount of information required to perform a particular task at a specified performance level. Hence it may provide an alternative technique to the classic Fitt's law measure of information [27]. For example, the spatial distribution of control behavior make it easier to determine the human equivalent control requirements (level of coordination, bandwidth) needed in difference regions of the planning space. These could also be used to develop new improved metrics to determine task difficulty and skill level. A sensitivity analysis based on the extremal fields would provide a rigorous way to investigate how pilot sensory/control uncertainties affect performance. For example, we can determine the cost penalty of inaccurate position estimation or poor velocity tracking. Such insights may provide additional clues on the principles at play in the biological control processes.

We plan to explore these ideas by conducting additional experiments. We plan to include more subjects to capture a wider range of skill levels. We also plan to incorporate uncertainties and disturbances in our tasks. This additional data will provide data to look further into the relationship between task difficulty, skill level, variability, and to further refine and validate our models of the human planning and control processes.

Ultimately, we hope that our framework will help better understand how the planning, control and sensing processes are implemented and organized in the brain to enable the pilot's unique control skills. Here, the map-like structure of the value function (being similar to other map-like representations used in the cognitive and neuro-sciences) may help provide a platform to better link psycho-motor data with biological models of the control and planning processes. This knowledge should eventually help us design more capable and versatile autonomous control algorithms as well as help design man-machine interfaces that provide a more natural link with the brain's processes involved in human control skills.

Acknowledgments: We are grateful to Nuri Kundak, Matt Rhinehart, and Hamid Mokhtarzadeh for their help during the data collection and to Zhaodan Kong for the computation of the optimal control trajectories.

REFERENCES

[1] D.T. McRuer and E.S. Krendel. "Mathematical Models of Human Pilot Behavior". Technical Report AGARD-AG-188, AGARD, 1974.

[2] V. Gavrilets, E. Frazzoli, B. Mettler, M. Piedmonte, and E. Feron. "Aggressive Maneuvering of Small Autonomous Helicopters: A Human-Centered Approach". *International Journal of Robotics Research*, pages 795–807, October 2001.

[3] V. Gavrilets, B. Mettler, and E. Feron. "Human-Inspired Control Logic for Automated Maneuvering of Miniature Helicopter". *AIAA Journal of Guidance, Control, and Dynamics*, 27(5), 2004.

[4] A. E. Bryson and Yu-Chi Ho. *Applied Optimal Control*. Taylor and Francis, Bristol, PA, 1975.

[5] John T. Betts. *Practical Methods for Optimal Control Using Nonlinear Programming*. SIAM, 2001.

[6] D.Q. Mayne, J.B. Rawling, C.V. Rao, and P.O.M. Scokaert. "Constrained model predictive control: Stability and optimality". *Automatica*, 36(6):789–814, 1987.

[7] A. Jadbabaie, J. Yu, and J. Hauser. "Unconstrained Receding-Horizon Control of Nonlinear Systems". *IEEE Transactions on Automatic Control*, pages 776–783, May 2001.

[8] A. Jadbabaie and J. Hauser. "Control of a thrust-vectored flying wing: a receding horizon–LPV approach". *International Journal of Robust and Nonlinear Control*, 12:869–896, 2002.

[9] James A. Primbs, Vesna Nevistic, and John C. Doyle. "A Receding Horizon Generalization of Pointwise Min-Norm Controllers". *IEEE Transactions on Automatic Control*, 45(5):898–909, May 2000.

[10] J. Bellingham, A. Richards, and J. How. Receding Horizon Control of Autonomous Aerial Vehicles. volume 5, pages 3741–3746. American Control Conference, May 2002.

[11] Y. Kuwata and J. How. "Three Dimensional Receding Horizon Control for UAVs". Number AIAA-2004-5144. American Institute of Aeronautics and Astronautics, 2004.

[12] B. Mettler and O. Toupet. Receding horizon trajectory planning with an environment based cost-to-go function. Seville, Spain, Dec. 2005. Proceedings of the joint ECC-CDC Conference.

[13] B. Mettler and Z. Kong. Receding horizon trajectory optimization with a receding horizon trajectory optimization with a finite-state value function approximation. American Control Conference, 2008.

[14] C. M. Harris and D. M. Wolpert. Signal-dependent noise determines motor planning. volume 394, pages 780–784. Nature, August 1998.

[15] D. Wolpert and Z. Ghahramani. Computational principles in movement neuroscience. *Nature Neuroscience*, 3:1212–1217, November 2000.

[16] E. Todorov and M. Jordan. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 3:1212–1217, November 2000.

[17] J. O'Keefe and L. Nadel. *The Hippocampus as a Cognitive Map*. Clarendon Press, Oxford, U.K.: Clarendon, 1978.

[18] T. Hafting, M. Fyhn, S. Molden, M.-B. Moser, and E. Moser. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(11):801–806, August.

[19] R. S. Sutton and A. G. Barto. *Reinforcement learning*. MIT Press, Cambridge, MA, 1998.

[20] A. Arleo and W. Gerstner. Spatial cognition and neuro-mimetic navigation: A model of hippocampal place cell activity, 2000.

[21] N. Kundak and B. Mettler. Experimental framework for evaluating autonomous guidance and control algorithms for agile aerial vehicles. Kos, Greece, July 2007. Proceedings of the European Control Conference (ECC).

[22] B. Mettler. *Identification Modeling and Characteristics of Miniature Rotorcraft*. Kluwer Academic Publishers, Boston, 2002.

[23] Eduardo D. Sontag. *Mathematical Control Theory: Deterministic Finite Dimensional Systems.* Springer, New York, NY, 1998.

[24] R. Bellman and S.E. Dreyfus. *Applied Dynamic Programming*. Princeton University Press, Princeton, New Jersey, 1962.

[25] A. Wachter and L.T. Biegler. On the implementation of a Primal-Dual Interior Point Filter Line Search Algorithm for Large-Scale Nonlinear Programming. Technical report, IBM T. J. Watson Research Center, USA, March 2004.

[26] A.P. Georgopoulos, J.F. Kalaska, and J.T. Massey. Spatial trajectories and reaction times of aimed movements: Effects of practice, uncertainty, and change in target location. *Journal of Neurophysiology*, 46(4):725–743.

[27] P. M. Fitts. The information capacity of the human motor system in controlling the amplitude of movements. *Journal of Experimental Psychology*, (47):381–391, 1954.